

## EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L7	0	(load adj3 balance adj manager) same resource	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:39
L6	2	(load adj3 balance adj manager) and L1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:39
L5	26	(distribut\$5) near5 (load adj3 balancer)and (policy rule) and ((resource server) adj3 (pool)) and L1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:38
L4	6	(distribut\$5) near5 (load adj3 balancer)and (policy rule) and ((resource server) adj3 (pool)) and "718"/\$.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:34
L3	2	(plural\$7 multiple) near5 (load adj3 balancer)and (policy rule) and ((resource server) adj3 (pool)) and "718"/\$.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:33
L2	14	(plural\$7 multiple) near5 (load adj3 balancer)and (policy rule) and ((resource server) adj3 (pool)) and L1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:32
S13 8	13	(plural\$7 multiple) near5 (load adj3 balancer)and (policy rule) and ((resource server) adj3 (pool))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/27 13:26



## EAST Search History

L1	56159	"709"/\$.ccls.	US-PGPUB; USPAT	OR	ON	2007/07/27 13:26
S13 9	58	("20010052006" "20010055317" "20020016856" "20020141401" "20030009561" "20030014525" "20030093496" "20030149690" "20040162901" "20050010754" "20060080446" "4766534" "5315708" "5396490" "5528761" "5553242" "5561807" "5745694" "5826082" "5937169" "5974518" "6091733" "6208650" "6212190" "6233615" "6304906" "6327622" "6341129" "6363421" "6411986" "6453360" "6480489" "6490281" "6526056" "6532487" "6549516" "6549961" "6606316" "6625650" "6628654" "6633560" "6650641" "6687222" "6704278" "6714985" "6724767" "6735169" "6742045" "6775692" "6781992" "6788704" "6836462" "6839811" "6891839" "6937606" "6973097" "7065086" "7072981").PN.	US-PGPUB; USPAT	OR	ON	2007/07/27 13:25
S14 2	1	10/777577	US-PGPUB; USPAT	OR	ON	2007/07/26 10:56
S14 0	27	S139 and (load near3 balanc\$5)	US-PGPUB; USPAT	OR	ON	2007/07/20 13:49
S1	0	((monitor\$4 Determin\$4) near10 (performanace)) and @ad<"19991126"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/20 13:42




[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "( ( multiple load balance&lt;in&gt;metadata ) ) &lt;and&gt; (pyr &gt;= 1950 &lt;and&gt; pyr &lt;= 20..."

e-mail

Your search matched 1 of 1621473 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)[New Search](#)

Modify Search

☐ Check to search only within this results setDisplay Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard

[Select All](#) [Deselect All](#)

1. **Negative-phase-sequence reduction with adjacent static reactive-power c**  
 Ledwich, G.; George, T.A.;  
[Electric Power Applications, IEE Proceedings-](#)  
 Volume 141, Issue 5, Sept. 1994 Page(s):259 - 263  
[AbstractPlus](#) | Full Text: [PDF\(320 KB\)](#) IET JNL

Indexed by  
 Inspec[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -



[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "( load balance manager&lt;in&gt;metadata )"

☒ e-mail

Your search matched 0 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)[New Search](#)

Modify Search

☐ Check to search only within this results setDisplay Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard

**No results were found.**

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.

Indexed by  
 Inspect[Help](#) [Contact Us](#) [Privacy & :](#) 

© Copyright 2006 IEEE -





USPTO

[Subscribe \(Full Service\)](#)   [Register \(Limited Service, Free\)](#)   [Login](#)

**Search:** ☒ The ACM Digital Library ☐ The Guide

+multiple +load +balancer resource



THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before March 2002

Terms used: **multiple load balancer resource**

Found 80 of 129,107

Sort results  
by

relevance

 Save results to a Binder

Try an Advanced Search

Try this search in The ACM Guide

Display results

expanded form ▼



## Search Tips

- Open results in a new window

Results 1 - 20 of 80

Result page: **1** 2 3 4 5 next

Relevance scale 

- # 1 Issues in the Design of Adaptive Middleware Load Balancing



Ossama Othman, Douglas C. Schmidt

August 2001 **ACM SIGPLAN Notices**, Proceedings of the ACM SIGPLAN workshop on Languages, compilers and tools for embedded systems LCTES '01, Proceedings of the 2001 ACM SIGPLAN workshop on Optimization of middleware and distributed systems OM '01, Volume 36 Issue 8

**Publisher:** ACM Press

Full text available:  pdf(239.32 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Load balancing middleware is used extensively to improve scalability and overall system throughput in distributed systems. Many load balancing middleware services are simplistic, however, since they are geared only for specific use-cases and environments. These limitations make it hard to use the same load balancing service for anything other than the distributed application it was designed for originally. This lack of generality forces continuous re-development of application-specific load balancing ...

**Keywords:** CORBA, load balancing, middleware, patterns, scalability

- ## 2 Heuristic methods for dynamic load balancing in a message-passing supercomputer

Jian Xu, Kai Hwang

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**  
**Supercomputing '90**

**Publisher:** IEEE Computer Society

Full text available:  pdf(1.04 MB)

Additional Information: full citation, abstract, references

In this paper, a new adaptive scheme is presented for dynamic load balancing on a message-passing multicomputer. The scheme is based on using easy-to-implement heuristics and variable threshold in migrating processes among the multicomputer nodes. It uses a distributed control over all processor nodes as coordinated by a host processor. Four heuristic methods for process migration are presented, which are distinguished by choosing different policies for process migration and threshold update. A ...

- ### 3 Cellular disco: resource management using virtual clusters on shared-memory




multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000



**ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3**Publisher:** ACM PressFull text available:  pdf(287.05 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a ...

**Keywords:** fault containment, resource management, scalable multiprocessors, virtual machines

#### 4 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5


**Publisher:** ACM PressFull text available:  pdf(1.93 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a syste ...

#### 5 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service

Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

**Publisher:** ACM PressFull text available:  pdf(2.52 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

**Keywords:** cluster, distributed systems, email, group membership protocol, load balancing, replication


#### 6 A dynamic load balancer for a parallel branch and bound algorithm

R. P. Ma, F-S. Tsung, M-H. Ma

January 1989 **Proceedings of the third conference on Hypercube concurrent computers and applications - Volume 2**


**Publisher:** ACM Press



Full text available:  pdf(756.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper presents a load-balancing scheme for a parallel branch and bound (PBB) algorithm utilizing a dynamic load balancer (DLB). The PBB algorithm is used to solve a resource scheduling problem on the Hypercube. The DLB is included in the PBB algorithm and is distributed to each node to balance the workload during run time. The DLB is evaluated by implementing two PBB algorithms, with DLB and without DLB. Results show that the PBB with DLB is 3.14 times faster than the algorithm without ...

## 7 [Dynamic adaptation to available resources for parallel computing in an autonomous network of workstations](#)

 Umit Rencuzogullari, Sandhya Dwardadas

June 2001 **ACM SIGPLAN Notices , Proceedings of the eighth ACM SIGPLAN symposium on Principles and practices of parallel programming PPOPP '01**, Volume 36 Issue 7

**Publisher:** ACM Press

Full text available:  pdf(344.17 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Networks of workstations (NOWs), which are generally composed of autonomous compute elements networked together, are an attractive parallel computing platform since they offer high performance at low cost. The autonomous nature of the environment, however, often results in inefficient utilization due to load imbalances caused by three primary factors: 1) unequal load (compute or communication) assignment to equally-powerful compute nodes, 2) unequal resources at compute nodes, and 3) multip ...

## 8 [Load balancing in a distributed processing system for high-energy physics\(UFMulti\)](#)


 Jagadeesh Kasaraneni, Theodore Johnson, Paul Avery

February 1995 **Proceedings of the 1995 ACM symposium on Applied computing SAC '95**

**Publisher:** ACM Press


Full text available:  pdf(623.69 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

## 9 [Distributed operating systems](#)

 Andrew S. Tanenbaum, Robbert Van Renesse


December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

**Publisher:** ACM Press

Full text available:  pdf(5.49 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

## 10 [Protecting web servers from distributed denial of service attacks](#)

 Frank Kargl, Joern Maier, Michael Weber

April 2001 **Proceedings of the 10th international conference on World Wide Web WWW '01**

**Publisher:** ACM Press

Full text available:  pdf(390.23 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



**Keywords:** DDoS, Linux, class based routing, distributed denial of service attacks, web server security

11 Dynamic load balancing for adaptive meshes using symmetric broadcast networks



Sajal K. Das, Daniel J. Harvey, Rupak Biswas

July 1998 **Proceedings of the 12th international conference on Supercomputing ICS '98**

**Publisher:** ACM Press

Full text available: pdf(1.01 MB) Additional Information: full citation, references, index terms

12 The effectiveness of multiple hardware contexts



Radhika Thekkath, Susan J. Eggers

November 1994 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the sixth international conference on Architectural support for programming languages and operating systems ASPLOS-VI, Volume 29 , 28 Issue 11 , 5**

**Publisher:** ACM Press

Full text available: pdf(1.26 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Multithreaded processors are used to tolerate long memory latencies. By executing threads loaded in multiple hardware contexts, an otherwise idle processor can keep busy, thus increasing its utilization. However, the larger size of a multi-thread working set can have a negative effect on cache conflict misses. In this paper we evaluate the two phenomena together, examining their combined effect on execution time. The usefulness of multiple hardware contexts depends on: program dat ...

13 Performance prediction of distributed load balancing on multicomputer systems



Ishfaq Ahmad, Arif Ghafoor, Kishan Mehrotra

August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing Supercomputing '91**

**Publisher:** ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, references, citations, index terms

14 Dataflow machine architecture



Arthur H. Veen

December 1986 **ACM Computing Surveys (CSUR), Volume 18 Issue 4**

**Publisher:** ACM Press

Full text available: pdf(3.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Dataflow machines are programmable computers of which the hardware is optimized for fine-grain data-driven parallel computation. The principles and complications of data-driven execution are explained, as well as the advantages and costs of fine-grain parallelism. A general model for a dataflow machine is presented and the major design options are discussed. Most dataflow machines described in the literature are surveyed on the basis of this model and its associated technology. F ...

15 Linux on Carrier Grade Web Servers



Ibrahim Haddad, Makan Pourzandi

April 2001 **Linux Journal**



**Publisher:** Specialized Systems Consultants, Inc.

Full text available:  [html\(20.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A great software solution for web traffic problems.


## 16 [Enabling dynamic content caching for database-driven web sites](#)



K. Selçuk Candan, Wen-Syan Li, Qiong Luo, Wang-Pin Hsiung, Divyakant Agrawal

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2

**Publisher:** ACM Press

Full text available:  [pdf\(319.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Web performance is a key differentiation among content providers. Snafus and slowdowns at major web sites demonstrate the difficulty that companies face trying to scale to a large amount of web traffic. One solution to this problem is to store web content at server-side and edge-caches for fast delivery to the end users. However, for many e-commerce sites, web pages are created dynamically based on the current state of business processes, represented in application servers and *databases*

**Keywords:** JDBC, application server, database driven web site, dynamic content caching, invalidation, web acceleration

## 17 [Scalable scheduling on a network of workstations](#)



Sanglu Lu, Li Xie

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

**Publisher:** ACM Press

Full text available:  [pdf\(675.97 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper presents an adaptive scalable scheduling model ASCS based on NOWs, which provides a dynamic load sharing facility by allowing tasks to scale up to utilize desired excess computing capacity depending upon available workstations on a NOW, and also scale down to maintain workstation autonomy. In this paper, a hierarchical scheduling policy is described: the inter-cluster scheduling provides a simple, fast and less-overhead load sharing among clusters in the system; and the intra-cluster ...

**Keywords:** coscheduling, load balancing, migration, scalable design, scheduling



## 18 [Scalable molecular dynamics for large biomolecular systems](#)



Robert K. Brunner, James C. Phillips, Laxmikant V. Kale

November 2000 **Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '00**

**Publisher:** IEEE Computer Society

Full text available:  [pdf\(211.89 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)  
 [Publisher Site](#)

We present an optimized parallelization scheme for molecular dynamics simulations of large biomolecular systems, implemented in the production-quality molecular dynamics program NAMD. With an object-based hybrid force and spatial decomposition scheme, and an aggressive measurement-based predictive load-balancing framework, we have attained speeds and speedups that are much higher than any reported in literature so far. The paper first summarizes the broad methodology we are pu ...

## 19



[Controlling application grain size on a network of workstations](#)





Bruce S. Siegel, Peter A. Steenkiste

December 1995 **Proceedings of the 1995 ACM/IEEE conference on Supercomputing (CDROM) - Volume 00 Supercomputing '95**

Publisher: ACM Press

Full text available: pdf(6.26 MB) Additional Information: full citation, abstract, references, citings, index terms  
html(2.56 KB)

An important challenge in the area of distributed computing is to automate the selection of the parameters that control the distributed computation. A performance-critical parameter is the grain size of the computation, i.e., the interval between successive synchronization points in the application. This parameter is hard to select since it depends both on compile time (loop structure and data dependences, computational complexity) and run time components (speed of compute nodes and network). On ...

**Keywords:** Automatic parallelization, grain size, network of workstations, dynamic load balancing

20 Cluster-based scalable network services

Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

Results 1 - 20 of 80

Result page: **1** 2 3 4 5 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player